arXiv:1511.08344v1 [cs.NI] 26 Nov 2015

# Scalable BGP Prefix Selection for Effective Inter-domain Traffic Engineering

Wenqin Shao, Luigi Iannone
and Jean-Louis Rougier
Telecom ParisTech
{firstname.lastname}@telecom-paristech.fr

François Devienne
and Mateusz Viste
Border 6
{firstname.lastname}@border6.com

***Abstract*—Inter-domain Traffic Engineering for multi-homed networks faces a scalability challenge, as the size of BGP routing table continue to grow. In this context, the choice of the best path must be made potentially for each destination prefix, requiring all available paths to be characterized (e.g., through measurements) and compared with each other. Fortunately, it is well-known that a few number of prefixes carry the larger part of the traffic. As a natural consequence, to engineer large volume of traffic only few prefixes need to be managed. Yet, traffic characteristics of a given prefix can greatly vary over time, and little is known on the dynamism of traffic at this aggregation level, including predicting the set of the most significant prefixes in the near future. Sophisticated prediction methods won't scale in such context.**

**In this paper, we study the relationship between prefix volume, stability, and predictability, based on recent traffic traces from nine different networks. Three simple and resource-efficient methods to select the prefixes associated with the most important foreseeable traffic volume are then proposed. Such proposed methods allow to select sets of prefixes with both excellent representativeness (volume coverage) and stability in time, for which the best routes are identified. The analysis carried out confirm the potential benefits of a route decision engine.**

## I. Introduction

BGP (Border Gateway Protocl) multi-homing has been a long-time common practice for a variety of large networks, such as enterprise networks, CDN (Content Delivery Networks), ISPs (Internet Service Provider), etc.[1] This solution is generally advocated for bringing enhanced connection reliability and improved traffic performance.

Yet, this introduces the task of actually taking advantage of the multiple paths that become available, when the network is connected to the Internet via multiple transit providers. In such context, one or more BGP border routers are needed for the connection with transit providers and also used for Traffic Engineering (TE) purposes. The best current practice in this context is to choose a best egress transit, in terms of certain optimization metrics, for each remote network (i.e., BGP prefix). These metrics represent path performance information, (e.g., Round-Trip Time, packet loss, etc), which are not part of path attributes BGP advertises. Fig. 1 provides a global picture of the tasks involved in performing TE in a multi-homed network. Several OAM (Operations, Administration, and Maintenance) function blocks are needed (cf., Fig. 1): (i) the performance of the paths through the different transits is measured; (ii) measures are then used by the "route decision engine", which selects the best egress transit for each BGP prefix and configure the BGP routers accordingly. One major challenge of these two functions is to cope with the growing Internet Routing Information base (RIB) [1]. However, it is well-known that most traffic is generally associated with only a fraction of all the possible BGP prefixes ( [2]–[5]), making it possible to focus only on such important prefixes.

In this work, we propose a framework to improve the scalability of fine-grained inter-domain TE in the context of BGP, by concentrating on the most important BGP prefixes. This is achieved by predicting which prefixes will

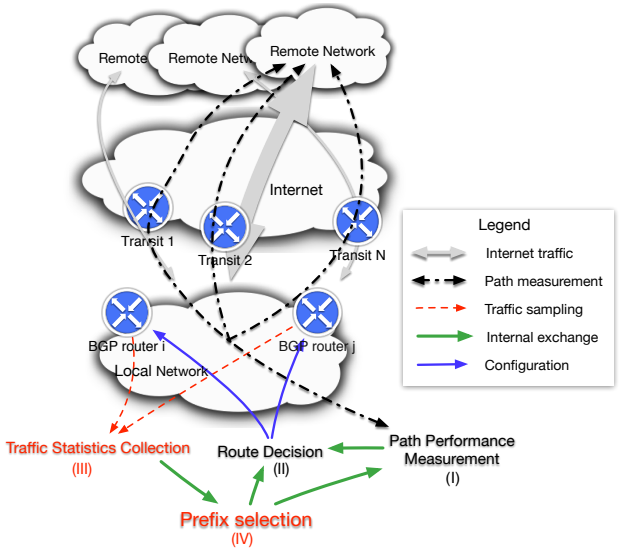


Fig. 1: Traffic Engineering with prefix selection.

[1]NAT-based multi-homing solutions, which do not require BGP, usually intended for smaller networks, are out of the scope of this paper.

correspond to the most important traffic volumes in the near future. To this end, two additional function blocks are needed (cf., Fig. 1): (iii) traffic analytics collection, which collects volume statistics; (iv) prefix selection process, which selects the set of the most important prefixes (i.e., the ones with the highest volume in the foreseeable future) and communicates the selected set to measurement and route decision function blocks.

Even if multi-homing has been used for more than a decade, there are still many challenges in fulfilling this goal. Well established Time Series Forecasting (TSF) models and artificial neural networks (ANN) have been used previously in traffic prediction [4], [6], [7]. These work targeted on highly aggregated inter-PoP (Point of Presence) traffic for off-line tasks, such as dimensioning.These models are computationally heavy, and require pre-treatment and tuning to fit well to each individual trace, which makes them less applicable in the context of inter-domain TE, where more than $10^5$ prefixes exist. Therefore, less complex prediction methods are needed. Furthermore, the dynamism of traffic variation makes prediction even more difficult. Traffic dynamism has been studied at a per-flow level. For instance, it has been shown by Walleriche et al. [8] that the bandwidth ranking of a 5-tuple flow can change drastically from one moment to another. It is however difficult to draw conclusions for the case of aggregated traffic (per BGP prefix). To our best knowledge, no study has given an in-depth investigation concerning the evolution in time of the traffic volume associated with BGP prefixes.

It is interesting to note that our work is similar to FIB (Forwarding Information Base) caching, which aims at installing only a small part of the routes, due to some software or hardware limitations. This problem has been extensively studied ( [5], [9]–[13]). We have compared our prefix selection methods to these works . We also believe that our study could be used for this problem as well.

The main contributions of this work are:

- A study of real traffic traces from networks of diverse profiles (ISPs, content providers and hosting providers) and from networks located in different countries (France, Germany, Poland, Spain, UK and USA).
- A through investigation on the Internet traffic dynamism across BGP prefixes was conducted on these traffic traces, which demonstrates diverse dynamism characteristics.
- The proposition of three scalable prediction methods. These methods are evaluated and compared to previous work. Our results show that the proposed schemes out-perform previous techniques both in term of performance and scalability.
- Finally, real measurements of the end-to-end path performance through each distinct transit provider for a set of prefixes is provided. The results confirmed older studies on multihoming in the fact that there could be significant differences among transit providers — highlighting, therefore, that the potential gain of inter-domain TE in the context of a multi-homed site can be substantial.

| | SA | SB | SC | SD | SE | SF | SG | SH | SI |
|---|---|---|---|---|---|---|---|---|---|
| **Type** | CP | ISP | HP | HP | CP | CP | ISP | CP | HP |
| **Vol.** | 133 | 528 | 6.7 | 1129 | 1871 | 5.1 | 0.2 | 29.9 | 6.2 |

TABLE I: Average traffic volume per hour (in GB) for the different measured networks

The rest of the article is organized as follows: in Sec. II, an in-depth study on the dynamism of traffic associated with BGP prefixes is given, based on which, three selection metrics are proposed and evaluated in Sec. III; Sec. IV evaluates the performance of different transit providers perceived by prefixes predictively selected and simulates the performance gain of a dynamic route selection algorithm based on latest RTT (Round-Trip Time) measurements toward selected prefixes; Sec. V presents previous related works; finally, we conclude our results and highlight the future directions in Sec. VI.

## II. Characters of Internet traffic over BGP prefixes

We base our study on working traffic traces collected from 9 networks of very different profiles listed in Table I. They are either Content Providers (CP), Internet Service Providers (ISP) or Hosting Providers (HP). Traffic trace from each network covers a time period of two entire weeks, from May 25th, 2015 to June 8th, 2015, except for SB and SD, for which the traces only cover the second week (starting June 1st, 2015). These traces were sampled from real traffic, similarly to previous works on FIB caching [11], [12]. It has been shown that the bias introduced by sampling is negligible in this kind of use cases.

Traffic volume for each prefixes are accumulated over 1 hour period. In comparison, studies concerning FIB caching [5], [12] use shorter time bins, from 1 second to 10 minutes, to capture instant changes of traffic variation. In our case, we argue that 1 hour is an appropriate update interval for prefix selection, as we are more interested in long term evolution.Changing the set of popular prefixes too frequently will make probing more difficult and might lead to frequent BGP route changes.

### *A. Traffic distribution over BGP prefixes*

As outlined earlier, this work is based on the fundamental assumption that traffic is concentrated on some popular traffic, as some previous work have shown [2], [3], [8]. For the sake of reliability, we quickly verify that our traffic traces demonstrate this property of concentration.

In Fig. 2a, the week volume share of each BGP prefix is represented. Prefixes are decreasingly sorted along the X-axis according to their cumulative volume fraction over the week. We observe that the weekly volume associated

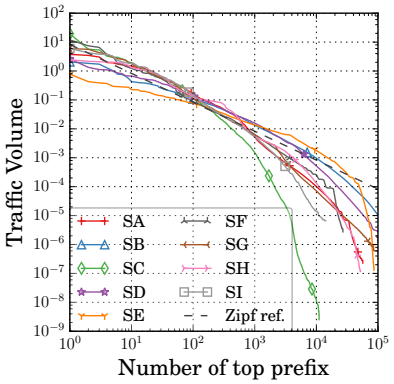


(a) Week volume share, log-log.

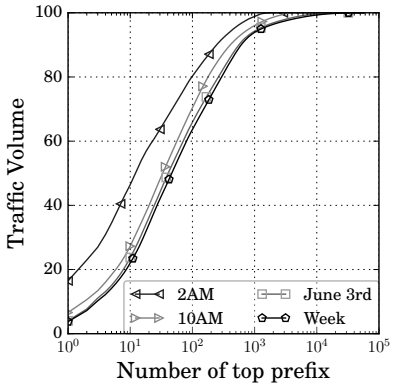


(b) CDF different time spans, SA.

Fig. 2: Traffic distribution among BGP prefixes.

with BGP prefixes can be approximately described by a reference Zipf's distribution with $N = 10^5, s = 1$ (dashed line)[2]. Fig. 2b compares the traffic distribution of SA over different time spans: volumes cumulated over one hour, at 2AM and 10AM, traffic cumulated over 24 hours (on June 3rd) and during the full week.

As expected, these graphs confirm that Internet traffic is highly concentrated on a few prefixes over multiple different periods of time, even after years of rapid RIB increase. However, Fig. 2b demonstrates that the level of traffic concentration over BGP prefixes is not stable but rather varies within a day, which leads to the study in the following section on traffic dynamism.

### *B. Dynamism of traffic over BGP prefixes*

In this section, we investigate the dynamism of traffic over BGP prefixes.

For each prefix $P$ ever active during a week, it has a time series $v(P) = \{v(P)_h\}_{h=1,\ldots,168}$ that stores its traffic volume each hour over the week. We define the Coefficient of Variation ($c_v$) of prefix $P$'s hourly volume series over the week as $c_v(P) = \frac{\delta(v(P))}{\mu(v(P))}$ , where $\delta$ denotes the standard deviation of the hourly volume series and $\mu$ is the mean hourly volume over the week. A large $c_v(P)$ value[3] indicates a large range of variation of traffic volumes w.r.t. its weekly average, and consequently more difficulty to anticipate the traffic volumes for this prefix $P$ [14].

Fig. 3 depicts the relation between the coefficient of variation and the "importance" of the prefix over a week. For each prefix $P$, the X-axis corresponds to its fraction of volume over the week, grouped in six bins: $[10^{-4}, 10^{-3})$, $\ldots$, to $[10, 100)$. In the Y-axis, we use a statistical representation, with mean, median, 25-75 percentile of

[2]Zipf's law defines that the $k^{th}$ most popular element among total $N$ elements has an occurrence share of $f(k, s, N) = \frac{1/k^s}{\sum_{n=1}^{N} 1/n^s}$.

[3]By construction, the maximum $c_v$ for a hourly volume series of 168 in length is $\sqrt{167}$, corresponding to the case where the prefix in question is active during only one single hour throughout the entire week.

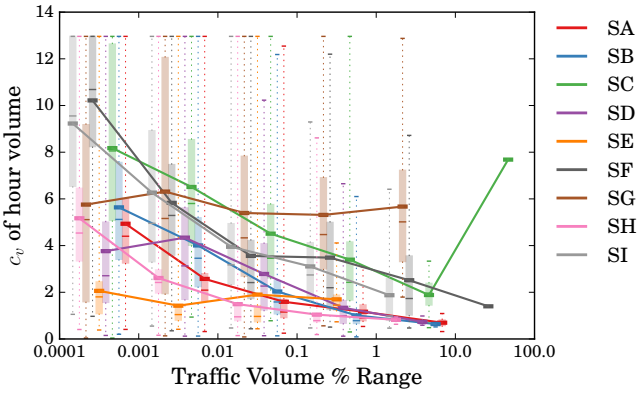


Fig. 3: Relation between $\{c_v(P)\}_P$ and fraction of weekly volume for all BGP prefixes $P$.

$\{c_v(P)\}_P$ (over all prefixes $P$ within the bin). We observe that, for all the networks except SG and SE, $c_v$ of prefixes that contribute to the biggest amount of traffic tends to be constrained in a narrower range around a smaller mean value. On the contrary, the prefixes which correspond to smaller weekly volume tend to have larger coefficient of variation. This is an interesting property, as we can capture a significant part of the overall traffic (represented by these stable prefixes) by picking the best prefixes based on their average hourly volume.

Next, we describe traffic dynamism from another perspective: For each hour $h$, we define the $core_h$ as the set containing top prefixes that represent 95% of total traffic. Table II lists the average number of prefixes included in the *core* and its percentage with regard to the average number of active prefixes each hour.

If a prefix is inside the *core* at a certain hour, it can be regarded important for bringing a significant amount of traffic. We thus defined the "*core* presence" for a prefix $P$ at each hour $h$ as: $cp(P)_h = 1$ when $P \in core_h$,and 0 otherwise. The *core* presence intensity over one week $I_{cp}(P, 168)$ is then defined as $I_{cp}(P, 168) = \frac{1}{168} \sum_{i=1}^{168} cp(P)_i$.

| Name | Avg. prefix # | Avg. prefix % | Max prefix # |
|---|---|---|---|
| SA | 629 | 17.87 | 1051 |
| SB | 5264 | 9.45 | 13934 |
| SC | 73 | 4.59 | 177 |
| SD | 2481 | 17.35 | 3757 |
| SE | 15501 | 53.61 | 20900 |
| SF | 377 | 30.73 | 772 |
| SG | 570 | 7.76 | 1766 |
| SH | 965 | 19.42 | 1731 |
| SI | 175 | 21.00 | 415 |

TABLE II: Core statistics.

| Network | Mean $BI$ | Max $BI$ | Max $\beta$ |
|---|---|---|---|
| SA | 14.61 | 37.79 | 7.44 |
| SB | 31.09 | 46.85 | 4.08 |
| SC | 40.57 | 145.07 | 145.05 |
| SD | 42.14 | 69.34 | 18.10 |
| SE | 20.91 | 44.30 | 20.17 |
| SF | 44.10 | 98.69 | 78.77 |
| SG | 51.21 | 125.41 | 102.05 |
| SH | 15.91 | 35.06 | 16.29 |
| SI | 38.59 | 85.47 | 56.56 |

TABLE III: Traffic burstiness.

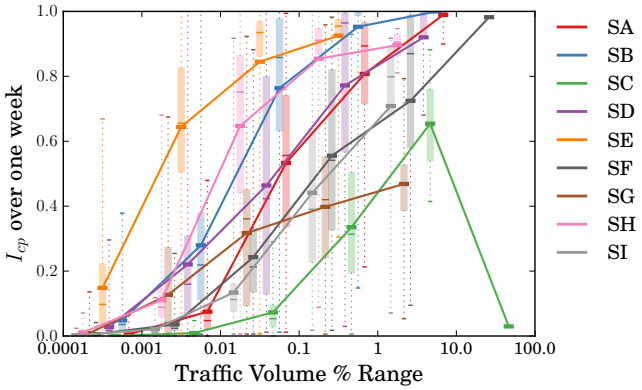


Fig. 4: Relation between $I_{cp}$ over the week and week volume fraction of BGP prefixes.

Fig. 4 plots the $I_{cp}$ over one week of each prefix (Y-axis) against its weekly volume fraction (in X-axis), using the same presentation as Fig. 3 described above. For all the networks, we can see that prefixes with bigger week volume share are less likely to have a low $I_{cp}$ over the week, i.e. they appear frequently in the *core*. We can conclude that, by focusing on prefixes that intensely appear in the *core* through the week, we will be able to capture a large part of the prefixes associated with important traffic volume over the week, and thus cover a large part of the overall traffic.

Finally, it is interesting to study the correlation between these two metrics: $c_v$ and $I_{cp}$. They are plotted together in Fig. 5 for two representative sites (SA and SC)[4], where each circle represents a prefix and its radius is proportional to the prefix in weekly volume share. The biggest circles mostly concentrate in the lower right corner of each sub-graph, which corresponds to the remarks made previously. However, some exceptions exist, especially on SC (but also SF, SG and SI), where we notice circles of big sizes having low $I_{cp}$ and relatively high $c_v$. These prefixes bring significant traffic volume share within a short duration, which makes predictive prefix selection difficult. Finally, it's not a surprise to see that the $c_v$ of prefixes with big week volume is, to a certain extent, inversely correlated to their $I_{cp}$.

[4]A comprehensive set of graphics is available on a research report [15]

### C. Quantitative index of traffic burstiness

In order to define traffic burstiness more formally and to systematically estimate this property, we define the index $\beta(P)_h$, for a prefix $P$ and a certain hour $h$ as: $\beta(P)_h = -\log(I_{cp}(P)) \times vp(P)_h$ if $I_{cp}(P) > 0$ and 0 otherwise, where $vp(P)_h$ is the hourly volume percentage of prefix $P$ at hour $h$. The logarithmic termon $I_{cp}$ aims at amplifying the volume contribution from prefixes with rare *core* presence, and attenuating the influence of prefixes being intensively in the *core*. A large $\beta$ value indicates that it is hard to predict the volume associated with this prefix while representing a significant hour volume.

In order to estimate the overall impact brought by bursty traffic at hour $h$, we sum up the $\beta(P)_h$ for each $P$ inside the *core* of that hour: more formally, $BI_h = \sum_{P \in core_h} \beta(P)_h$. For all the networks, we estimate their traffic burstiness with the mean and maximum value of $BI$ series over the week. The results are given in Table III. The maximum $\beta$ over all the week is also given for the sake of clarity.

We found that this simple metric is able to capture the burstiness of a given site. The existence of prefixes representing significant volumes but with fairly low $I_{cp}$ , such as for the site SC in Fig. 5, correspond to the biggest maximum $\beta$ values. On the contrary, a small value of $BI$ (around 30 or lower) is arrived for sites such as SA, with more predictable volume-significant prefixes.

## III. Predictive prefix selection

Based on the previous observations, several measures are possible for predicting the"importance" of a prefix. They are enumerated below.

*Mean Volume:* With this metric, we predict that at hour $h+1$, the volume importance of prefix $P$ is indicated by it's mean hourly volume over the last $L$ hours, $MV(P,L)_{h+1} = 1/L \times \sum_{i=h-L+1}^{h} v(P)_i$. It is based on the observations from Fig. 3, that top prefixes ranked over the week tend to have smaller hourly volume variation around their mean volume.

*Core Presence Intensity:* The prediction could use $I_{cp}(P,L)_{h+1} = 1/L \times \sum_{i=h-L+1}^{h} cp(P)_i$, i.e. the *core* presence intensity of the prefix over the last $L$ hours. It derives from the observation made from Fig. 4, that high

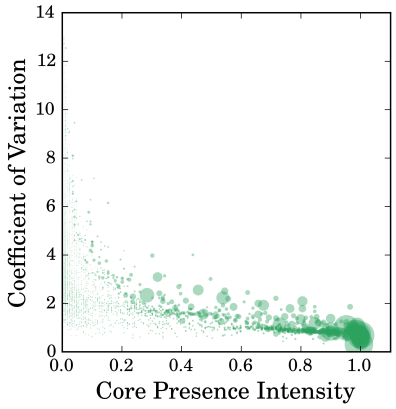


(a) SA, 56684 prefixes

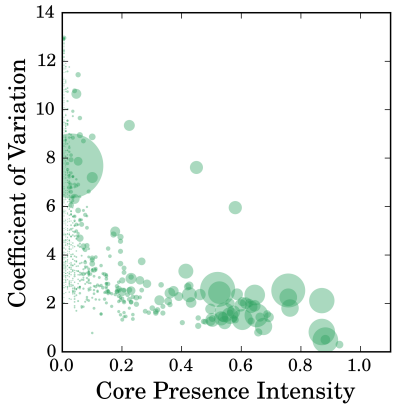


(b) SC, 11065 prefixes

Fig. 5: Relation between $I_{cp}$ and $c_v$

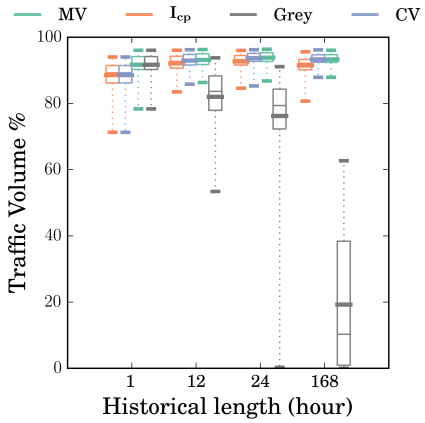


(a) SA

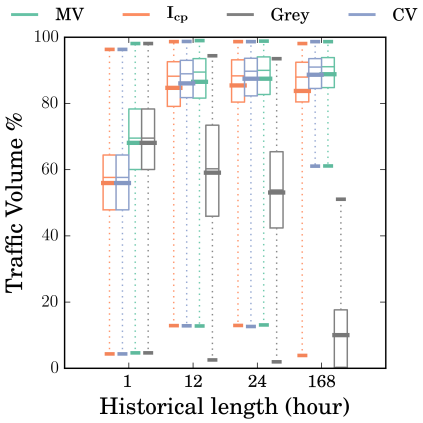


(b) SC

Fig. 6: Hourly volume fraction covered by the selected prefixe set

ranked prefixes by their week volume are more likely to have intense *core* presence.

*Core Volume:* Finally, the prediction could be based on $CV(P,L)_{h+1} = 1/L \times \sum_{i=h-L+1}^{h} cp(P)_i \times v(P)_i$, a combination of both $MV$ and $I_{cp}$. $CV$ has the potential to be more resource thrifty compared to $MV$, as it is be calculated only for those prefixes which ever appeared in the *core* over the last $L$ hours — while $MV$ is computed for all active prefixes. From our observations, the *core* size actually represents 5% to 50% of all active prefixes.

In previous work by Zhange et al. [12] on FIB caching, a grey differential model $GM(1,1)$ [16] is employed to predict which BGP prefixes will represent the biggest packet counts. For the sake of comparison, we implemented the $GM(1,1)$ model as well in our work.

We fix the selection size to the maximum *core* size over the week. We can then evaluate the quality of the prefix set selected with the different proposed metrics, or with the $GM(1,1)$ model.

Fig. 6 illustrates the hourly volume coverage by the four methods in the form of boxplot, representing the minimum, maximum, 25th and 75th percentile, medium and mean values. Among proposed metrics, we find that the $CV$ is very close to $MV$ in terms of volume coverage, proving that it is a good approximation of the later, while being more resource economical.

Basing solely on the last 1 hour records, all methods yield already a mean volume coverage $> 80\%$ on SA, SB, SE and SH, which implies a strong continuity on prefix volumes between two consecutive hours. On SC and SG, however, using records of last 168 hours, i.e. a week, offers much better minimum volume coverage than shorter records. This is due to the fact that at certain hour, SC and SG undergo a great amount of bursty traffic (as observed previously, e.g. in Table III and Fig. 5). By increasing the historical length, the selection metrics are able to have better visibility into the past and capture some of these bursty prefixes — finally improving the minimum coverage. This gain in minimum volume coverage by using long historical records can actually be observed on all

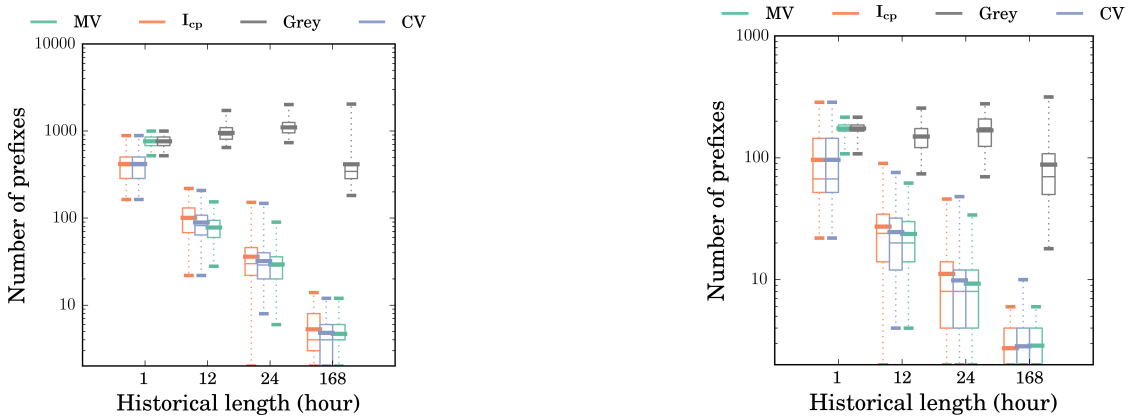


(a) SA (b) SC

Fig. 7: Hourly churn of the selected prefix set

networks, between 1 hour and 12 hours, also between 12 hour and 24 hour. However from 24 hour to 168, this gain doesn't necessary happen on all sites, which is due to the fact that the total volume brought by some highly bursty prefix is diluted by the long time span using $MV$ and $CV$ metrics. In order to capture them, a larger selection set size is need, which inevitably includes more prefixes of few significance.

On the other hand, the hourly volume coverage by grey model drops as we increase the historical records length and is in general much worse than the metrics proposed in this work. In order to understand the underlying reason, we used $GM(1,1)$ model described above to dynamically predict the total hourly volume of all prefixes, which is normally much more regular and smoother than the volume series of individual prefixes. It sorts out that the grey model reacts to variations in a delayed manner when $L$ being big and extravagantly overreacts to sudden changes when $L$ is small [15].

Fig. 7 gives the results concerning the churn of the selection prefix set (with a boxplot representation). The churn is defined as the difference (i.e. number of new and deleted prefixes) between the new predicted set and the previous one. A high churn could pose operational problems: when probing is considered, for instance, a stable set is clearly more adapted. Sarrar et al. [5] also argued that small prefix churn is very important in network architecture with decoupled forwarding and control planes, such as SDN (Software Defined Networking), as it leads to lower communication overhead.

As expected, a clear drop of the churn value can be seen when the historical length increases — as opposed to what happens with the grey model. In that sense, using long historical records can be a wise choice in practical uses. Furthermore, for networks with relatively few bursty traffic, e.g. SA, the mean volume coverage with last 168 hour records is extremely close to that with last 24 hour records, as seen in Fig. 6a. Finally, for networks with highly bursty traffic, SC, using long records has the potential to obviously improve worst-case volume coverage.

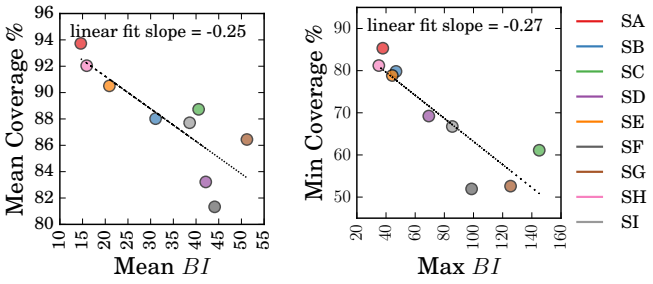


(a) Average level (b) Worst case

Fig. 8: The relationship between burstiness index $BI$ and traffic volume coverage of selected prefix using $CV$ metric.

In the purpose of lowering churn, Sarrar et al. [5] proposed selecting top prefixes over time bins of different lengths (ranging from 1 second to 10 minute in their FIB-caching environment). In our context, we found that the difference in mean volume coverage using record lengths larger than 1 hour is marginal, thus little gain can be expected from this method.

Finally, in Fig. 8 the mean/minimum coverage achieved with $CV$ metric is showed as a function of the $BI$ index. We can see that the mean (resp. minimum) coverage is inversely proportional to the mean (resp. maximum) $BI$ index. As expected, this graphic highlight the difficulty to cover a large trafic volume for networks with more bursty traffic. However, the quasi-linear curve obtained proves that $BI$ is a very meaningful metric to identify sites with bursty trafic. For large $BI$ values, it can be worthy of choosing a larger prefix set size (which was previously fixed to the weekly maximum core size), if possible.

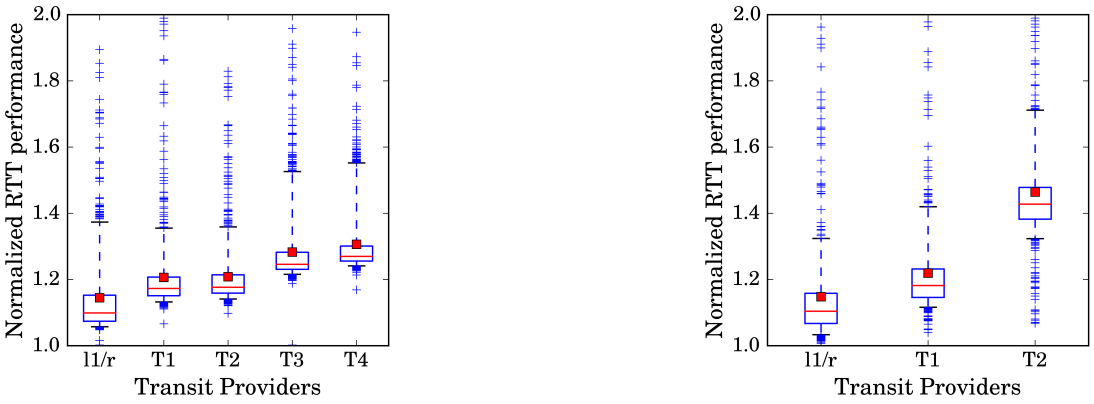


(a) SA (b) SC

Fig. 9: Normalized RTT performance with active probing.

## IV. RTT probing and transit performance evaluation

We then study the potential performance that our multi-homing optimization system can gain. We first present the method used to evaluate the transit performance in terms of RTT (Round-Trip Time) perceived by selected prefixes.

TCP SYN scanning was employed to measure RTT between selected prefixes and the local network. The probe traffic toward a selected prefix was steered on all available transit providers, by means of explicit routing. For a pair of selected prefix and transit provider, a probe is scheduled at 240 second interval in average with 30% randomization left or right. The probe data covers the entire week starting from June 1st.

Our performance evaluation method follows the same principle as the one proposed by Akella et al. [17]: The computation is done as :

$$NP_{t_i}^{Tx} = \frac{1}{|SP|} \sum_{P \in SP} M_{t_i}^{Tx}(P) / \min_{Tj \in T} M_{t_i}^{Tj}(P) \qquad (1)$$

where $NP_{t_i}^{Tx}$ is the normalized RTT performance for transit provider $Tx$ at probe tickle $t_i$, and $M_{t_i}^{Tx}(P)$ denotes the RTT measured toward selected prefix $P$.

This measure is first normalized over the best measure among all available transit providers $T$ at the same probe tickle. Then we average this normalized RTT over all prefixes inside the selected prefix set $SP$. If a transit provider offers the smallest RTT to all selected prefixes, it should have a normalized RTT performance equaling to 1.

Fig. 9 gives the results of transit performance evaluation for SA, in a "boxplot" format, where the intervals represent $5^{th}$ and $95^{th}$ percentiles (values below and beyond, marked by a + symbol, can be regarded as outliers). Along the X-axis, available transit providers are aligned increasingly according to their mean normalized RTT performance (marked by a square).

$l1/r$ denotes a virtual transit provider simulated by a dynamic route selection algorithm inspired from [18]. The method selects the transit provider that provides the smallest RTT in last round of probing for each selected prefix. This choice is based on the hypothesis that RTT of a path demonstrates temporal locality and thus is closely related to its most recent measurement. It is a basic algorithm used for the sake of illustration only.

The number of prefixes probed on each network and traffic volume represented are given in Fig. 6a.

We observe that the performance differences among different transit providers is particularly clear, in particular on SC. This confirms that multi-homing can still provide significant performance improvement nowadays. However this gain in performance is not inherently given in the context of BGP. Even for transit provider that offers the best best mean $NP$, there exist moments where its performance deviates far above 1, which means that traffic toward some selected prefixes are suffering from RTTs much larger than on other available transit providers. This implies that a network will not have optimal performance if it uses transit providers in a static and indistinguishable manner for each individual prefixes. Therefore a dynamic route selection method (such as $l1/r$) that dynamically select the best outgoing next-hop for each individual selected prefix is needed. $l1/r$ selects the transit that provides the smallest RTT in last probing round. In the case that measurement data is not available, it chooses randomly. For all networks except SG, the simulated algorithm out-performs all physical transit providers. On SB, the overall RTT is 20% below the mean RTT of the best available transit provider. Still the $NP$ value of this virtual provider varise within a wide range, which calls for further investigation on the characters of RTT variation in time.

## V. Related work

Some previous works [3], [18], [19] acknowledged the importance of preforming inter-domain TE only for destinations that matter. However, no general solution was

proposed to predictively select the BGP prefixes with important traffic volume. Some other works leveraged the skewed distribution of Internet traffic to downsize FIB [5], [9]–[13], where understanding of traffic dynamism on much shorter time scales, e.g. seconds and minutes, are required. In our work, we used traffic traces of coarser time resolution over longer period of time — conditions more adapted to the complicated operations associated required for inter-domain TE.

Zhang et al. [12] assumed that the stability and popularity of Internet traffic is positively correlated without verification. Papagiannaki et al. [20] showed that this correlation doesn't necessarily exist for 5-tuple flows on 5 minute interval.

Akella et al. [17] quantified the performance gain from multi-homing using traces from a large CDN network more than ten years ago. In a latter work [18], they evaluated a dynamic route selection system on a testbed, but with only 100 destinations.

## VI. Conclusion and future directions

The present paper tackled the problem of controlling the majority of data traffic via a small subset of BGP prefixes, by exploiting the uneven Internet traffic distribution. The main goal being the simplification of management procedures associated with inter-domain Traffic Engineering, in particular for multi-homed sites. One of the challenges in addressing such problem was to select in a scalable manner the prefixes that will carry most traffic volume in the forthcoming time.

We analyzed real traffic measurements from nine different sites (located in five different countries) to understand the distribution of traffic volume associated with different BGP prefixes, as well as its variation in time. We observed how the most important prefixes (representing largest volume over a week) are generally stable in time, with small hourly variations around their mean. Based on our observations, we proposed three simple metrics (also easy to compute) to proactively select prefixes with important foreseeable traffic volume. We demonstrated that the metrics we proposed lead to better volume coverage compared to the existing solutions. Furthermore, we measured RTT performance through different transit providers and simulated a route decision engine on the selected prefix set. Through such analysis we showed that with such dynamic route selection the overall RTT performance of the network can be improved by 20% compared to the best available transit provider.

In order to achieve better prefix selection methods, we have shown that capturing bursty prefixes is the key point. Interesting directions to explore in order to improve the already good results we achieved is to group prefixes by their activity profiles. For each group, selection method could be adapted to the its traffic dynamism. When dealing with prefixes with regular volume patterns and small hourly variation, the simple metrics proposed in the work perform extremely well and may suffice. Nevertheless, for bursty prefixes, we might need a more sophisticate model that extracts additional activity features, like for instance long term periodicity. The use of the burstiness index $\beta$, shown to be very expressive, in prefix classification is also worthy of further work.

## References


[1] BGP Routing Table Analysis Reports. [Online]. Available: http://bgp.potaroo.net

[2] W. Fang and L. Peterson, "Inter-AS traffic patterns and their implications," in *GLOBECOM '99*, vol. 3. IEEE, 1999, pp. 1859–1868.

[3] N. Feamster, J. Borkenhagen, and J. Rexford, "Guidelines for interdomain traffic engineering," *ACM SIGCOMM CCR*, 2003.

[4] K. Papagiannaki, N. Taft, Z.-L. Zhang, and C. Diot, "Long-term forecasting of internet backbone traffic." *IEEE transactions on neural networks*, vol. 16, no. 5, pp. 1110–24, Sep. 2005.

[5] N. Sarrar, S. Uhlig, A. Feldmann, R. Sherwood, and X. Huang, "Leveraging Zipf's law for traffic offloading," *ACM SIGCOMM CCR*, vol. 42, no. 1, p. 16, Jan. 2012.

[6] P. Cortez, M. Rio, M. Rocha, and P. Sousa, "Internet Traffic Forecasting using Neural Networks," in *IJCNN '06*. IEEE, 2006, pp. 2635–2642.

[7] T. Otoshi, Y. Ohsita, M. Murata, Y. Takahashi, K. Ishibashi, and K. Shiomoto, "Traffic Prediction for Dynamic Traffic Engineering Considering Traffic Variation," *GLOBECOM '13*, pp. 1592–1598, Dec. 2013.

[8] J. J. Wallerich and A. Feldmann, "Capturing the variability of internet flows across time," *INFOCOM '06*, pp. 1–6, 2006.

[9] L. Iannone and O. Bonaventure, "On the cost of caching locator/ID mappings," *CoNEXT '07*, p. 1, 2007.

[10] H. Ballani, P. Francis, T. Cao, and J. Wang, "Making routers last longer with ViAggre," *NSDI '09*, pp. 453–466, 2009.

[11] C. Kim, M. Caesar, A. Gerber, and J. Rexford, "Revisiting route caching: The world should be flat," *Lecture Notes in Computer Science*, vol. 5448, pp. 3–12, 2009.

[12] W. Zhang, J. Bi, J. Wu, and B. Zhang, "Catching popular prefixes at AS border routers with a prediction based method," *Computer Networks*, vol. 56, no. 4, pp. 1486–1502, Mar. 2012.

[13] Y. Liu, V. Lehman, and L. Wang, "Efficient FIB caching using minimal non-overlapping prefixes," *Computer Networks*, vol. 83, pp. 85–99, 2015.

[14] Q. He, C. Dovrolis, and M. Ammar, "On the predictability of large transfer TCP throughput," in *SIGCOMM '05*, vol. 35, no. 4. New York, New York, USA: ACM Press, Aug. 2005, p. 145.

[15] W. Shao, L. Iannone, J.-L. Rougier, F. Devienne, and V. Mateusz. Research report on prefix selection for BGP TE. [Online]. Available: http://biblio.telecom-paristech.fr/cgi-bin/download.cgi?id=15740

[16] D. Julong, "Introduction to Grey System Theory," *The Journal of Grey System1*, vol. 1, pp. 1–24, 1989.

[17] A. Akella, B. Maggs, S. Seshan, A. Shaikh, and R. Sitaraman, "A measurement-based analysis of multihoming," in *SIGCOMM '03*. New York, New York, USA: ACM Press, 2003, p. 353.

[18] A. Akella, B. Maggs, S. Seshan, and A. Shaikh, "On the Performance Benefits of Multihoming Route Control," *IEEE/ACM Transactions on Networking*, vol. 16, no. 1, pp. 91–104, Feb. 2008.

[19] D. K. Goldenberg, L. Qiuy, H. Xie, Y. R. Yang, and Y. Zhang, "Optimizing cost and performance for multihoming," *ACM SIGCOMM CCR*, vol. 34, no. 4, p. 79, Oct. 2004.

[20] K. Papagiannaki, N. Taft, and C. Diot, "Impact of Flow Dynamics on Traffic Engineering Design Principles," in *INFOCOM '04*, vol. 4, 2004, pp. 2295–2306.